1

Game Theory: A Lexicon for Strategic Interaction

> High-rationality solution concepts in game theory can emerge in a world populated by low-rationality agents.
>
> *Young (1998)*

1.1 Introduction

This chapter defines and illustrates the fundamental tools of game theory: extensive form games, normal form games, and Nash equilibria. The reader is asked to take the importance of the Nash equilibrium concept on faith for the time being, since its justification requires the development of conceptual tools—evolutionary stability and dynamics—that appear only several chapters down the road. The chapter begins, however, with a carefully developed example—Big Monkey and Little Monkey—that showcases a major contribution of game theory to understanding social interactions: insisting on absolute clarity concerning the rules of the game and its informational assumptions. In the process, we touch on such engaging issues as credible and incredible threats and the value of precommitment, thus setting the tone for the next several chapters.

I have found that class time is best used by presenting the example in detail and the formal definitions relatively sketchily, relying on the problem set (§1.6) to drive home the subtleties of the formalities.

1.2 Big Monkey and Little Monkey

Big Monkey and Little Monkey normally eat fruit and berries from small ground-level bushes. But an important part of their diet is warifruit, which dangle from the extreme tip of a lofty branch of the waritree. A waritree only occasionally produces fruit, and then only one fruit per tree. To get the warifruit, at least one of the monkeys must climb the tree, creep part

VISIT…

way out the branch on which the fruit is growing, and shake vigorously until the fruit comes loose and falls to the ground. Careful energy measurements show that a warifruit is worth 10 Kc (kilocalories) of energy, the cost of running up the tree, shaking the fruit loose, and running back down to the ground costs 2 Kc for Big Monkey, but is negligible for Little Monkey, who is much smaller. Moreover, if both monkeys climb the tree, shake the fruit loose, then climb down the tree and eat the fruit, Big Monkey gets 7 Kc and Little Monkey gets only 3 Kc, since Big Monkey hogs most of it; if only Big Monkey climbs the tree, while Little Monkey waits on the ground for the fruit to fall, Big Monkey gets 6 Kc and Little Monkey gets 4 Kc (Little Monkey eats some before Big Monkey gets back down from the tree); if only Little Monkey climbs the tree, Big Monkey gets 9 Kc and Little Monkey gets 1 Kc (most of the food is gone by the time Little Monkey gets there).

What will Big Monkey and Little Monkey do if each wants to maximize net energy gain? We may safely put aside for the moment some of bigger questions, such as: (a) How do we know monkeys maximize anything? (b) How do monkeys come to know the costs and the benefits of the various actions? (c) Are monkeys smart enough to find an optimal solution? (d) Who cares about monkeys, anyway? We will come back to these larger issues later. We simply assume that monkeys do maximize, they do know the costs and benefits, they can find the optimal solution, they cannot enforce cooperative agreements, and, finally, that we will learn something interesting by solving this problem.

There is one more matter that must be resolved: Who decides first what to do, Big Monkey or Little Monkey? There are three possibilities: (a) Big Monkey decides first; (b) Little Monkey decides first; (c) both monkeys decide simultaneously. We will go through the three cases in turn.

Assuming Big Monkey decides first, we get the situation depicted in Fig. 1.1. We call a figure like this a *game tree*, and we call the game it defines an extensive form game. At the top of the game tree is the *root node* (the little dot labeled "Big Monkey") with two *branches*, labeled w (wait) and c (climb). This means Big Monkey gets to choose and can go either left (w) or right (c). This brings us to the two nodes labeled "Little Monkey," in each of which Little Monkey can wait (w) or climb (c).

Note that while Big Monkey has only two strategies, Little Monkey actually has four:

a. Climb no matter what Big Monkey does (cc).
b. Wait no matter what Big Monkey does (ww).

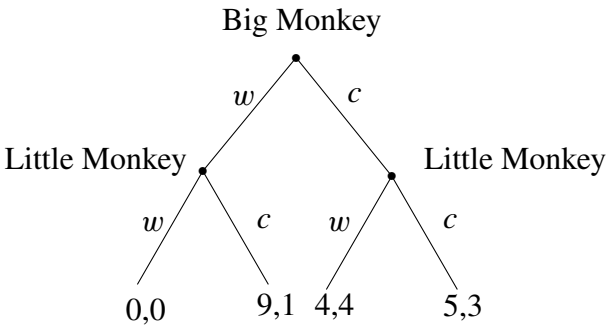

Figure 1.1. Big Monkey and Little Monkey: Big Monkey chooses first.

c. Do the same thing Big Monkey does (wc).
d. Do the opposite of what Big Monkey does (cw).

We call a move taken by a player (one of the monkeys) at a node an *action*, and we call a series of actions that fully define the behavior of a player a *strategy*—actually a *pure strategy*, in contrast to "mixed" and "behavioral" strategies, which we will discuss later, that involve randomizing. Thus, Big Monkey has two strategies, each of which is simply an action, while Little Monkey has four strategies, each of which is two actions—one to be used when Little Monkey is on the left, and one when Little Monkey is on the right.

At the bottom of the game tree are four nodes, which we variously call *leaf* or *terminal nodes*. At each terminal node is the payoff to the two players, Big Monkey (player 1) first and Little Monkey (player 2) second, if they choose the strategies that take them to that particular leaf. You should check that the payoffs correspond to our description above. For instance, at the leftmost leaf when both wait, with neither Monkey expending or ingesting energy, the payoff is (0,0). At the rightmost leaf both climb the tree, costing Big Monkey 2 Kc, after which Big Monkey gets 7 Kc and Little Monkey gets 3 Kc. Their net payoffs are thus (5,3). And similarly for the other two leaves.

How should Big Monkey decide what to do? Well, Big Monkey should figure out how Little Monkey will react to each of Big Monkey's two choices, w and c. If Big Monkey chooses w, then Little Monkey will choose c, since this pays 1 Kc as opposed to 0 Kc. Thus, Big Monkey gets 9 Kc by moving left. If Big Monkey chooses c, Little Monkey will choose w, since this pays 4 Kc as opposed to 3 Kc for choosing c. Thus Big Monkey gets 4 Kc for

choosing c, as opposed to 9 Kc for choosing w. We now have answered Big Monkey's problem: choose w.

What about Little Monkey? Well, Little Monkey must certainly choose c on the left node, but what should Little Monkey choose on the right node? Of course it doesn't really matter, since Little Monkey will never *be* at the right node. However, we must specify not only what a player does "along the path of play" (in this case the left branch of the tree), but at *all possible nodes on the game tree*. This is because we can only say for sure that Big Monkey is choosing a best response to Little Monkey if we know what Little Monkey does, and conversely. If Little Monkey makes a wrong choice at the right node, in some games (though not this one) Big Monkey would do better by playing c. In short, Little Monkey must choose one of the four strategies listed above. Clearly, Little Monkey should choose cw (do the opposite of Big Monkey), since this maximizes Little Monkey's payoff no matter what Big Monkey does.

Conclusion: the only reasonable solution to this game is for Big Monkey to wait on the ground, and Little Monkey to do the opposite of what Big Monkey does. Their payoffs are (9,1). We call this a Nash equilibrium (named after John Nash, who invented the concept in about 1950). A Nash equilibrium in a two-player game is a pair of strategies, each of which is a *best response* to the other; i.e., each gives the player using it the highest possible payoff, given the other player's strategy.

There is another way to depict this game, called its *strategic form* or *normal form*. It is common to use both representations and to switch back and forth between them, according to convenience. The normal form corresponding to Fig. 1.1 is in Fig. 1.2. In this example we array strategies of player 1 (Big Monkey) in rows, and the strategies of player 2 (Little Monkey) in columns. Each entry in the resulting matrix represents the payoffs to the two players if they choose the corresponding strategies.

		Little Monkey			
		cc	cw	wc	ww
Big Monkey	w	9,1	9,1	0,0	0,0
	c	5,3	4,4	5,3	4,4

Figure 1.2. Normal form of Big Monkey and Little Monkey when Big Monkey moves first.

We solve the normal form of the game by trying to pick out a row and a

column such that the payoff to their intersection is the highest possible for player 1 down the column, and the highest possible for player 2 across the row (there may be more than one such pair). Note that (w, cw) is indeed a Nash equilibrium of the normal form game, because 9 is better than 4 for Big Monkey down the cw column, and 1 is the best Little Monkey can do across the w row.

Can we find any other Nash equilibria to this game? Clearly (w, cc) is also a Nash equilibrium, since w is a best reply to cc and conversely. But the (w, cc) equilibrium has the drawback that if Big Monkey should happen to make a mistake and play c, Little Monkey gets only 3, whereas with cw, Little Monkey gets 4. We say cc is *weakly dominated* by cw.

But what if Little Monkey plays ww? Then Big Monkey should play c, and it is clear that ww is a best response to c. So this gives us another Nash equilibrium, (c, ww), in which Little Monkey does much better, getting 4 instead of 1, and Big Monkey does much worse, getting 4 instead of 9. Why did we not see this Nash equilibrium in our analysis of the extensive form game? The reason is that (c, ww) involves Little Monkey making an *incredible threat*. "I don't care what you do, says Little Monkey—I'm waiting here on the ground—no matter what." The threat is "incredible" because Big Monkey knows that if he plays w, then when it is Little Monkey's turn to carry out the threat to play w, Little Monkey will not in fact do so, simply because 1 is better than 0.[1] We say a Nash equilibrium of an extensive form game is *subgame perfect* if, at any point in the game tree, the play dictated by the Nash equilibrium *remains* a Nash equilibrium of the subgame (we give formal definitions and more examples in chapter 5). The strategy (c, ww) is not subgame perfect because in the subgame beginning with Little Monkey's choice of w on the left of Fig. 1.1 is not a best response. Nice try, anyway, Little Monkey!

But what if Little Monkey gets to choose first? Perhaps now Little Monkey can force a better split than getting 1 compared to Big Monkey's 9. This is the extensive form game (Fig. 1.3). We now call Little Monkey player 1 and Big Monkey player 2. Now Big Monkey has four strategies (the strategies that belonged to Little Monkey in the previous version of the game) and Little Monkey only has two (the ones that belonged to Big Monkey before). Little Monkey notices that Big Monkey's best response to w is c, and Big Monkey's best response to c is w. Since Little Monkey gets 4 in the first

[1]This argument fails if the monkeys can condition their behavior in one day on their behavior in previous days (see chapter 6). We assume the monkeys cannot do this.

case and only 1 in the second, Little Monkey chooses w. Big Monkey's best choice is then cw, and the payoffs are (4,4). Note that *by going first, Little Monkey is able to precommit to a strategy that is an incredible threat when going second*.

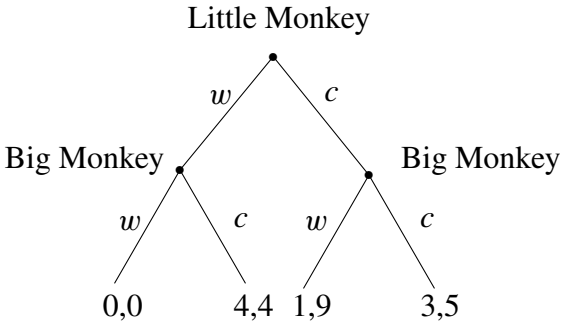

Figure 1.3. Big Monkey and Little Monkey: Little Monkey chooses first.

The normal form for the case when Little Monkey goes first is illustrated in Fig. 1.4. Again we find the two Nash equilibria (w, cc) and (w, cw), and again we find another Nash equilibrium not evident at first glance from the game tree: now it is Big Monkey who has an incredible threat, by playing ww, to which Little Monkey's best response is c.

		Big Monkey			
		cc	cw	wc	ww
Little Monkey	w	4,4	4,4	0,0	0,0
	c	3,5	1,9	3,5	1,9

Figure 1.4. Normal form of Big Monkey and Little Monkey game when Little Monkey moves first.

The final possibility is that the monkeys choose simultaneously or, equivalently, each monkey chooses an action without seeing what the other monkey chooses. In this case, each monkey has two options: climb the tree (c), or wait on the ground (w). We then get the situation in Fig. 1.5. Notice the new element in the game tree: the dotted line connecting the two places where Little Monkey chooses. This is called an *information set*. Roughly speaking, an information set is a set of nodes at which (a) the same player chooses, and (b) the player choosing does not know which particular node

represents the actual choice node. Note also that we could just as well interchange Big Monkey and Little Monkey in the diagram, reversing their payoffs at the terminal nodes, of course. This illustrates an important point: there may be more than one extensive form game representing the same real strategic situation.

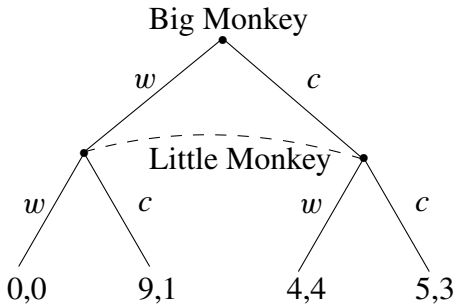

Figure 1.5. Big Monkey and Little Monkey choose simultaneously.

Even though there are fewer strategies in this game, it is hard to see what an equilibrium might be by looking at the game tree. This is because what Little Monkey does cannot depend on which choice Big Monkey makes, since Little Monkey does not see Big Monkey's choice. So let's look at the normal form game, in Fig. 1.6. From this figure, it is easy to see that both (w, c) and (c, w) are Nash equilibria, the first obviously favoring Big Monkey and the second favoring Little Monkey. In fact, there is a third Nash equilibrium that is more difficult to pick out. In this equilibrium Big Monkey randomizes by choosing c and w with probability $1/2$, and Little Monkey does the same. This is called a *mixed strategy Nash equilibrium*; you will learn how to find and analyze it in chapter 4. In this equilibrium Big Monkey has payoff 4.5 and Little Monkey has payoff 2. The reason for this meager total payoff is that with probability $1/4$, both wait and get zero reward, and sometimes both climb the tree!

		Little Monkey: c	Little Monkey: w
Big Monkey	c	5,3	4,4
	w	9,1	0,0

Figure 1.6. Big Monkey and Little Monkey: normal form in the simultaneous move case.

1.3 The Extensive Form Game

Now that you have seen an example of one, actually defining a game should not be daunting. The formal definition of an extensive form game is *notationally challenged* it requires much mathematic notation to express something that is conceptually simple. We will forego some of this in favor of a verbal description; but where the notation cannot be avoided without causing confusion, we will not hesitate to use it. An *extensive form game* $\mathcal{G}$ consists of a number of *players*, a *game tree*, and a set of *payoffs*. A game tree consists of a number of *nodes* connected by *branches*. Each branch connects a *head node* to a distinct *tail node* If b is a branch of the game tree, we denote the head node of b by b^h, and the tail node of b by b^t.

A *path* from node a to node a' in the game tree is a sequence of moves starting at a and ending at a'.[2] If there is a path from node a to a', we say a is an *ancestor* of a', and a' is a *successor* to a. We call k the *length* of the path. If a path from a to a' has length one, we call a the *parent* of a', and a' is a *child* of a. We require that the game tree have one node r, called the *root node*, that has no parent, and a set T of nodes called *terminal nodes* or *leaf nodes*, that have no children. We associate with each terminal node $t \in T$, and each player i, a *payoff* $\pi_i(t) \in \mathbf{R}$ ($\mathbf{R}$ is the set of real numbers). We say the game is *finite* if it has a finite number of nodes. We assume all games are finite, unless otherwise stated.

We also require that the graph of $\mathcal{G}$ have the following *tree property*. There must be *exactly one path* from the root node to any given terminal node in the game tree. One obvious but important implication of this property is that *every node except the root node has exactly one parent*.

Players relate to the game tree as follows. Each nonterminal node is assigned to a player who moves at that node. Each branch from a node represents a particular *action* that the player assigned to that node can take there, and hence determines either a terminal node or the next point of play in the game—the particular child node to be visited next.[3]

If a stochastic event occurs at a node a (for instance, the weather is Good or Bad, or your partner is Nice or Nasty), we assign the fictitious player

[2]Technically, a path is a sequence $b_1, \ldots, b_k$ of branches such that $b_1^h = a$, $b_i^t = b_{i+1}^h$ for $i = 1, \ldots k-1$, and $b_k^t = a'$; i.e., the path starts at a, the tail of each branch is the head of the next branch, and the path ends at a'.

[3]Thus if $\mathbf{p} = (b_1, \ldots, b_k)$ is a path from a to a', then starting from a, if the actions associated with the b_j are taken by the various players, the game moves to a'.

"Nature" to that node, the actions Nature takes representing the possible outcomes of the stochastic event, and we attach a *probability* to each branch of which a is the head node, representing the probability that Nature chooses that branch (we assume all such probabilities are strictly positive).

The tree property thus means that there is a *unique* sequence of moves by the players (including Nature) leading from the root node to any specific node of the game tree, and for any two nodes, there is *at most one* sequence of player moves leading from the first to the second.

A player may know the exact node in the game tree when it is his turn to move (e.g., the first two cases in Big Monkey and Little Monkey, above), but he may know only that he is at one of several possible nodes. This is the situation Little Monkey faces in the simultaneous choice case (Fig. 1.6). We call such a collection of nodes an *information set*. For a set of nodes to form an information set, the same player must be assigned to move at each of the nodes in the set and have the same array of possible actions at each node.

For some purposes we also require that if two nodes a and a' are in the same information set for a player, the moves that player made up to a and a' must be the same. This criterion is called *perfect recall*, because if a player never forgets his moves, he cannot make two different choices that subsequently land him in the same information set.[4] We will assume perfect recall unless otherwise stated in this book, but whenever a result depends upon perfect recall, we will mention the fact.

Suppose each player $i = 1, \ldots, n$ chooses strategy s_i. We call $s = (s_1, \ldots, s_n)$ a *strategy profile* for the game, and we define the *payoff* to player i, given strategy profile s, as follows. If there are no moves by Nature, then s determines a unique path through the game tree, and hence a unique terminal node $t \in T$. The payoff $\pi_i(s)$ to player i under strategy profile s is then defined to be simply $\pi_i(t)$. We postpone the case where Nature moves to §11.2.1, since the definition depends on the so-called *expected utility principle*.

[4]Another way to describe perfect recall is to note that the information sets $\mathcal{N}_i$ for player i are the nodes of a graph in which the children of an information set $\nu \in \mathcal{N}_i$ are the $\nu' \in \mathcal{N}_i$ that can be reached by one move of player i, plus some combination of moves of the other players and Nature. Perfect recall means that this graph has the tree property.

1.4 The Normal Form Game

The *strategic form* or *normal form* game consists of a number of players, a set of strategies for each of the players, and a payoff function that associates a payoff to each player with a choice of strategies by each player. More formally, the n-player normal form game consists of

a. A set of *players* $i = 1, \ldots, n$.
b. A set S_i of *strategies* for player $i = 1, \ldots, n$. We call $s = (s_1, \ldots, s_n)$, where $s_i \in S_i$ for $i = 1, \ldots, n$, a *strategy profile* for the game.[5]
c. A function $\pi_i : S \to \mathbf{R}$ for player $i = 1, \ldots, n$, where S is the set of strategy profiles, so $\pi_i(s)$ is player i's payoff when strategy profile s is chosen.

Two extensive form games are said to be *equivalent* if they correspond to the same normal form game, except perhaps for the labeling of the actions and the naming of the players. But given an extensive form game, how exactly do we form the corresponding normal form game? First, the players in the normal form are the same as the players in the extensive form. Second, for each player i, let S_i be the set of strategies of that player, each strategy consisting of a choice of an action at each information set where i moves. Finally, the payoff functions are given by equation (3.2). If there are only two players and a finite number of strategies, we can write the payoff function in the form of a matrix, as in Fig. 1.2.

1.5 Nash Equilibrium

The concept of a Nash equilibrium of a game is formulated most easily in terms of the normal form. Suppose the game has n players, with strategy sets S_i and payoff functions $\pi_i : S \to \mathbf{R}$, for $i = 1, \ldots, n$ where S is the set of strategy profiles. We use the following very useful notation. If $s \in S$ and $t_i \in S_i$, we write

$$(s_{-i}, t_i) = (t_i, s_{-i}) = \begin{cases} (t_1, s_2, \ldots, s_n) & \text{if } i = 1 \\ (s_1, \ldots, s_{i-1}, t_i, s_{i+1}, \ldots, s_n) & \text{if } 1 < i < n \\ (s_1, \ldots, s_{n-1}, t_n) & \text{if } i = n \end{cases}.$$

In other words, (s_{-i}, t_i) is the strategy profile obtained by replacing s_i with t_i for player i.

[5]Technically, these are *pure strategies*, because later we will consider *mixed strategies* that are probabilistic combinations of pure strategies.

We say a strategy profile $s^* = (s_1^*, \ldots, s_n^*) \in S$ is a *Nash equilibrium* if, for every player $i = 1, \ldots, n$, and every $s_i \in S_i$, we have $\pi_i(s^*) \geq \pi_i(s_{-i}^*, s_i)$; i.e., choosing s_i^* is at least as good for player i as choosing any other s_i given that the other players choose s_{-i}^*. Note that in a Nash equilibrium, the strategy of each player is a *best response* to the strategies chosen by all the other players. Finally, notice that a player could have responses that are *equally good* as the one chosen in the Nash equilibrium—there just cannot be a strategy that is strictly better.

The Nash equilibrium concept is important because in many cases we can accurately (or reasonably accurately) predict how people will play a game by assuming they will choose strategies that implement a Nash equilibrium. It will also turn out that, in dynamic games that model an evolutionary process whereby successful strategies drive out unsuccessful ones over time, stationary states are, with some obvious and uninteresting exceptions, Nash equilibria. Moreover, dynamic evolutionary systems often either converge to a Nash equilibrium (we then say it is *asymptotically stable*) or follow closed orbits around a Nash equilibrium (we then say it is *neutrally stable*). We will also see that many Nash equilibria that seem implausible intuitively are actually *unstable* equilibria of an evolutionary process, so we would not expect to see them in the real world. Where people appear to deviate systematically from implementing Nash equilibria and in no sense appear ignorant or "irrational" for doing so, we will usually find that we have misspecified the game they are playing or the payoffs we attribute to them.

By contrast with our "evolutionary" approach, the older tradition of classical game theory justifies the Nash equilibrium concept by reference to the concept of *common knowledge*. We say that a fact p is "common knowledge" if all players know p, all players know that all players know p, all players know that all players know that all players know p, and so on ad infinitum. One can then show that if it is common knowledge that all players are rational (i.e., choose best responses), then under appropriate conditions they will choose a Nash equilibrium (see Tan and da Costa Welang 1988 and the references therein). In his famous entry in *The New Palgrave*, for instance, Aumann (1987b) reflects the received wisdom in asserting, "The common knowledge assumption underlies all of game theory [otherwise] the analysis is incoherent."[6]

However often repeated, such assertions are simply false. Common knowl-

[6]Actually, more recent research shows that even within classical game theory this assumption can often be considerably weakened (Aumann and Brandenburger 1995).

edge is a *sufficient* but not a *necessary* condition for agents to choose strategies forming a Nash equilibrium (and then only under quite restrictive conditions). Moreover, it is a ridiculously implausible assumption even for humans. Finally, the successful application of game theory to biology (see, for instance, §4.15 and §4.16), and the successful artificial life simulation of games (§4.11, §5.11, §7.14 and §13.8), show that such strong assumptions are unnecessary. Hence, unless otherwise stated (as for instance in §3.21), in this book we will use the term "common knowledge" to mean that all players know the fact in question, with no additional levels of recursion implied.

1.6 Reviewing the Terminology

The following exercises require that you understand fully and completely the concepts developed above.

a. Define the following terms: path, successor, ancestor, parent, child, root node, terminal node, information set, action.
b. Any extensive form game can be translated into a normal form game in a natural manner. Use the Big Monkey and Little Monkey example to describe how this is done, and why.
c. Conversely, any normal form game can be translated into an "equivalent" extensive form game. How is this done? In what sense are the two "equivalent"? Using the Big Monkey and Little Monkey example, show that, unlike the reverse direction, this procedure is not unique.
d. What is a *best response* in a normal form game? What is a Nash Equilibrium?
e. Write out the conditions for a pair of strategies to form a Nash equilibrium of a two-player normal form game.
f. A *graph* is any set of nodes and branches, a branch being a line segment connecting together two nodes. A *directed graph* is a graph in which, for each branch, one of the nodes is designated the *head* node and the other the *tail* node. So a game tree is a special sort of directed graph. What properties must we add to a directed graph to get a game tree for an extensive form game?